



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ :	A1	(11) International Publication Number: WO 94/19888
H04L 12/24, 12/26, H04M 3/24, G06F 11/22		(43) International Publication Date: 1 September 1994 (01.09.94)
<p>(21) International Application Number: PCT/US94/02056</p> <p>(22) International Filing Date: 25 February 1994 (25.02.94)</p> <p>(30) Priority Data: 08/023,972 26 February 1993 (26.02.93) US</p> <p>(71) Applicant: CABLETRON SYSTEMS, INC. [US/US]; 35 Industrial Way, Rochester, NH 03867-5005 (US).</p> <p>(72) Inventor: LEWIS, Lundy; 480 Greenville Road, Mason, NH 03048 (US).</p> <p>(74) Agent: MORRIS, James, H.; Wolf, Greenfield & Sacks, 600 Atlantic Avenue, Boston, MA 02210 (US).</p>		(81) Designated States: AU, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>
(54) Title: METHOD AND APPARATUS FOR RESOLVING FAULTS IN COMMUNICATIONS NETWORKS		
<p>(57) Abstract</p> <p>An improved method and apparatus of resolving faults in a communications network. The preferred system uses a trouble ticket data structure to describe communications network faults. Completed trouble tickets are stored in a library and when an outstanding trouble ticket is received, the system uses at least one determinator to correlate the outstanding communications network fault to data fields in the set of data fields of the trouble ticket data structure to determine which completed trouble tickets in the library are relevant to the outstanding communications network fault. The system retrieves a set of completed trouble tickets from the library that are similar to the outstanding trouble ticket and uses at least a portion of the resolution from at least one completed trouble ticket to provide a resolution of the outstanding trouble ticket. The determinators may be macros, rules, a decision tree derived from an information theoretic induction algorithm and/or a neural network memory derived from a neural network learning algorithm. The system may adapt the resolution from a retrieved trouble ticket to provide the resolution using null adaptation, parameterized adaptation, abstraction/respecialization adaptation, or critic-based adaptation techniques.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australin	GE	Georgin	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KR	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

-1-

METHOD AND APPARATUS FOR RESOLVING FAULTS
IN COMMUNICATIONS NETWORKS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to communications networks, and in particular to a method and apparatus for resolving faults in such networks. Within this disclosure, the term "communications network" is used to refer to any type of digital communications system, of which a computer-based, local area network or a computer-based, wide area network are examples.

2. Discussion of the Related Art

All communications networks experience faults during network operation. Faults, as used in this disclosure, may include a failure of hardware portions of the communications network, such as workstations or peripheral devices and failure of software portions of the network, such as software application programs and data management programs. In small stable homogeneous communications networks (i.e., those in which all of the equipment is provided by the same vendor and the network configuration does not change), management and repair of network faults is relatively straightforward. However, as a communications network becomes increasingly large and heterogeneous (i.e., those in which different types of equipment are connected together over large areas, such as an entire country), fault management becomes more difficult.

One of the ways to improve fault management in large communications networks is to use a so-called "trouble-ticketing" system. This system provides a number of tools that can be used by network users, administrators, and repair and maintenance personnel. The basic data structure, a "trouble ticket," has a number of fields in which a user can enter data describing the parameters of an observed

-2-

network fault. A trouble ticket filled out by a user may then be transmitted by, for example, an electronic mail system to maintenance and repair personnel. A trouble ticket describing a current network fault that needs to be acted on is called an "outstanding trouble ticket". When the network fault has been corrected, the solution to the problem, typically called a "resolution" is entered into an appropriate data field in the trouble ticket. When a network fault has been resolved, the trouble ticket is said to be completed. The system provides for storage of completed trouble tickets in a memory and thus a library of such tickets is created, allowing users, administrators, and maintenance and repair personnel to refer to these stored completed trouble tickets for assistance in determining solutions to new network faults.

The trouble-ticketing system thus provides a convenient, structured way of managing fault resolution and for storing solutions to network faults in a manner that allows this stored body of knowledge to be accessed and applied to outstanding communications network faults. An example of a trouble-ticketing system is the ACTION REQUEST SYSTEM, developed by Remedy Corporation, Mountainview, California, and sold by Cabletron Systems, Inc., Rochester, New Hampshire.

A structured trouble-ticketing system, however, does not provide a complete solution to the fault management problem. For time-critical network services, the downtime that elapses from the observation of a network fault, the submission of a trouble ticket, to the completion of the trouble ticket can be expensive. Downtime can be reduced by providing a communication link between a network fault detection system and a trouble-ticketing system. The communication link allows fault information collected by the fault detection system to be transmitted to the trouble-ticketing system in the form of an automatically-generated and filled out trouble ticket. The trouble-ticketing system then manages

-3-

communication and workflow among the network administrator, support staff, and end-users, in the normal manner to resolve the outstanding trouble ticket.

Although this solution allows trouble tickets to reach the fault management system and appropriate maintenance and repair personnel more quickly, it does not reduce the time necessary to resolve an outstanding fault. A maintenance and repair person is still required to research and resolve the outstanding fault. This is not only time-consuming, but expensive as well.

To reduce the time in which faults are resolved, artificial intelligence systems may be used to assist in resolving the outstanding trouble ticket. In existing systems that make use of artificial intelligence in this manner, fault resolution expertise is represented using a rule-based reasoning (hereinafter RBR) method.

A typical RBR system includes a working memory, a rule-base, and a control procedure. The working memory typically contains a representation of characteristics of the network, including topological and state information. The rule-base represents knowledge about what operations should be performed when the network malfunctions. If the network enters an undesirable state, the control procedure selects those rules that are applicable to the current situation. Of the rules that are applicable, a predetermined control strategy selects a rule to be executed. A rule can perform tests on the network, query a database, provide commands to a network configuration management system, or invoke another expert system. Using results obtained after executing a rule, the system updates the working memory by asserting, modifying, or removing working memory elements. The RBR system continues in this cycle until a desirable state in the working memory representing a desirable state of the network is achieved. Examples of RBR systems for network management

-4-

may be seen in Expert Systems Applications in Integrated Network Management, edited by E. Erickson, L. Ericson, D. Minoli and published by Archtech House, Inc., 1989.

Constructing an RBR fault resolution system requires defining a description language that appropriately and completely represents networking conditions (the "domain"), extracting expertise from persons with expertise in the network ("domain experts") and/or trouble-shooting documents, and representing the expertise in the RBR format. This procedure requires several iterations of a so-called "consult/implement/test" cycle in order to achieve a correct system. In the consult/implement/test cycle, an expert is interviewed to determine his or her fault resolution methodology, the methodology is implemented in a rule or rules that the system can process, and the rules are tested. If the conditions or domain in which the RBR system operates remains relatively stable, once a correct system is achieved, minimal maintenance is required. However, if the system is used to resolve faults in unpredictable or rapidly changing domains, two problems typically occur. First, the RBR system suffers from the problem of "brittleness". Brittleness means that the system fails when it is presented with a novel problem for which it has no applicable rules. A cause of system brittleness is that the system cannot adapt existing knowledge to a novel situation or cannot gain new information from novel experiences to apply in the future. The second problem is commonly known as a "knowledge acquisition bottleneck". The knowledge acquisition bottleneck occurs when a knowledge engineer tries to manually modify the rule-base by devising special rules and control procedures in order to deal with changes, new parameters, or other unforeseen situations. As a result of these modifications, the RBR system typically becomes unwieldy, unpredictable, and unmaintainable. Furthermore, if the domain in which the RBR system operates is a rapidly changing one, the system can become obsolete in a relatively short period of time.

-5-

Therefore, an object of the present invention is to provide a method and apparatus for resolving faults in communications networks that learns from prior fault resolution scenarios and offers solutions to novel network faults based on past resolution scenarios.

Another object of the present invention is to provide a method and apparatus that applies so-called case-base reasoning (hereinafter "CBR") to fault management and resolution in communications networks.

Still another object of the present invention is to provide a method and apparatus for automatically resolving faults in communications networks.

SUMMARY OF THE INVENTION

The present invention overcomes the disadvantages of the prior art by providing an improved method and apparatus for resolving faults in a communications network. The system broadly includes acquiring information about an outstanding communications network fault, retrieving from a database relevant, previously stored sets of data regarding prior communications network faults and resolutions of those prior communications network faults, and using at least a portion of at least one of the previously stored resolutions to provide a resolution of the outstanding communications network fault. More specifically, the system utilizes a trouble ticket data structure capable of describing information concerning a communications network fault and using the trouble ticket data structure to create a trouble ticket that describes a specific network fault, the trouble ticket having a set of data fields capable of storing information regarding the network fault and a data field capable of storing a resolution of the specific network fault. A trouble ticket library capable of storing at least one trouble ticket is provided. When an outstanding trouble ticket is received, the system uses at least one determinator to correlate the outstanding fault to data fields in the set

-6-

of data fields of the trouble ticket data structure and retrieves at least one trouble ticket from the library that is similar to the outstanding trouble ticket as a function of the at least one determinator. The determinator is a mechanism that allows the system to focus on and retrieve relevant trouble tickets and reduces the likelihood of the system retrieving irrelevant trouble tickets. The determinator identifies and correlates network faults with data fields in the trouble ticket that store data representing network characteristics or operating parameters and that are relevant to determining a resolution of the network fault. Determinators may be described as macros, rules, a decision tree determined by an information theoretic induction algorithm and/or as a neural network memory determined by a neural network learning algorithm. Thereafter, the system uses at least a portion of the resolution from at least one stored ticket to provide a resolution of the outstanding ticket.

In one embodiment of the invention, the system adapts the resolution from at least one similar retrieved trouble ticket to provide a resolution of the outstanding trouble ticket using parameterized adaptation, abstraction/respecialization adaptation, and/or critic-based adaptation techniques. The system can receive and complete outstanding trouble tickets automatically or in conjunction with user input.

The features and advantages of the present invention will be more readily understood and apparent from the following detailed description of the invention, which should be read in conjunction with the accompanying drawings, and from the claims which are appended at the end of the detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

In the drawings, which are incorporated herein by reference and in which like elements have been given like reference characters,

-7-

FIG. 1 is a block diagram of a fault management and resolution system of the present invention;

FIG. 2 is a block diagram of a trouble-ticketing system that may be used in the system of FIG. 1;

FIG. 3 is an example of a trouble ticket to be processed by the system of FIG. 1;

FIG. 4 is a flow chart illustrating the processing steps carried out by the fault management and resolution system of FIG. 1;

FIG. 5 is a flow chart illustrating the processing steps carried out by the retrieve module of FIG. 4;

FIG. 6 is a flow chart illustrating the processing steps carried out by the adapt module of FIG. 4; and

FIG. 6A is a flow chart illustrating processing steps carried out by the propose module of FIG. 4.

FIG. 7 illustrates a pseudo-code model that may be used to implement the processing functions illustrated in FIGS. 4-6A.

DETAILED DESCRIPTION

For purposes of illustration only and not to limit generality, the present invention will now be explained with reference to its use in management and resolution of faults occurring in a typical computer-based local area network. However, one skilled in the art will recognize that the present invention is applicable to other types of communications networks.

FIG. 1 is an overall block diagram of a fault management and resolution system 10 which incorporates the present invention. The system includes a network monitoring system 12 that is coupled via communications link 14 to a local area network 8. Network monitoring system 12 is coupled, via communication link 16 to fault processing system 18. By "communications link" is meant any type of software/hardware data transmission/reception medium.

-8-

Network monitoring system 12 may be the SPECTRUM™ local area network monitoring system manufactured by Cabletron Systems, Inc., Rochester, New Hampshire. The SPECTRUM™ system is described in copending and commonly owned application Serial No. 07/583,509 entitled NETWORK MANAGEMENT SYSTEM USING MODEL-BASED INTELLIGENCE filed September 17, 1990; Serial No. 07/789,000 entitled NETWORK MANAGEMENT SYSTEM USING STATUS SUPPRESSION TO ISOLATE NETWORK FAULTS filed November 7, 1991; Serial No. 07/790,408 entitled NETWORK MANAGEMENT SYSTEM USING INTERCONNECTED HIERARCHIES TO REPRESENT DIFFERENT DIMENSIONS filed November 7, 1991; Serial No. 07/788,936 entitled NETWORK MANAGEMENT SYSTEM USING MULTIFUNCTION ICONS filed November 7, 1991; Serial No. 07/797,121 entitled METHOD AND APPARATUS FOR MONITORING THE STATUS OF NON-POLABLE DEVICES IN A COMPUTER NETWORK filed November 22, 1991, which are hereby incorporated by reference in their entirety.

The network monitoring system 12 includes a configuration management module 20 and fault detection module 22. Configuration management module 20 provides for monitoring and control of certain aspects of network operation, such as uploading or downloading software, allocating memory, and network organization. Fault detection module 22 monitors local area network 8 via communications link 14, configuration management module 20 and communications link 24 to detect any undesirable network conditions that indicate a fault has occurred. If a network fault is detected, fault detection module 22 may automatically gather and transmit appropriate fault information via communications link 16 to fault processing system 18.

Fault processing system 18 includes a gateway or interface 26 that receives fault information from fault detection module 22 and applies any necessary communication protocol with the fault detection module 22 and automatically, in one embodiment of the invention, gateway 26 generates a trouble ticket that may be processed by fault

-9-

management and resolution system 28. Gateway 26 may be the ARS GATEWAY manufactured by Cabletron Systems, Inc., Rochester, New Hampshire. The so-processed trouble ticket is transmitted to fault management and resolution system 28 via communications link 30.

Fault management and resolution system 28 includes a fault resolution system 32 coupled to a fault management or trouble-ticketing system 34. Trouble-ticketing system 34 may be the aforementioned ACTION REQUEST SYSTEM. As will be explained in greater detail, fault resolution system 32 processes a received trouble ticket, retrieves relevant trouble tickets that may be useful in resolving the particular network fault from trouble-ticketing system 34, adapts (in one embodiment of the invention) resolutions from the retrieved trouble tickets to address the network fault of the outstanding trouble ticket, and processes the resolution automatically or under user control to correct the communications network fault. The fault resolution may be transmitted via communications link 36 to configuration management module 20, for use on network 8.

Fault processing system 18 also includes a user interface module 38 coupled to fault resolution system 32 via communications link 40 and trouble-ticketing system 34 via communications link 42. User interface module 38 allows a user to edit and control proposed fault resolutions generated by fault resolution system 32 using keyboard 44. Proposed resolutions and system status may be viewed by the user on display 46 or permanently recorded by printer 48.

The functions carried out by fault processing system 18 may be implemented by programming a general purpose computer as will be described in greater detail. The general purpose computer may be the network file server, a client workstation attached to the network, or a workstation dedicated to use by maintenance and repair personnel for network fault resolution.

-10-

FIG. 2 is a block diagram of various modules included in trouble-ticketing system 34. A trouble ticket memory 50 is used to store completed trouble tickets and provides a library of fault resolution scenarios that may be called up by the fault resolution system 32 for dealing with an outstanding trouble ticket. An administrator module 52, allows the system administrator to configure the trouble ticket format for various system users. A notification module 54 provides automatic notification to a pre-selected maintenance and repair person upon receipt of a trouble ticket. A user module 56 allows a user to access trouble ticket memory 50 in order to view a trouble ticket and allows entry or modification of certain data in the trouble ticket such as repair status by authorized personnel. A communications link 58 connects the various modules together.

FIG. 3 illustrates the format of a typical trouble ticket automatically generated by gateway 26 to be processed by fault resolution system 32 and trouble-ticketing system 34. Trouble ticket 60 includes a plurality of fields 62A through 62R. These fields enable a user manually or network monitoring system 12 automatically to fill out the trouble ticket so that a maintenance and repair person may obtain the information necessary to resolve the problem. As noted, trouble ticket 60 was automatically generated by gateway 26 using data from network monitoring system 12 and therefore field 62G contains SPECTRUM™, the name of the network monitoring system as the submitter's name. The fault detection module 22 automatically provides the data such as the submitter's name, the creation date, the alarm condition, etc. for entry into the appropriate field of the trouble ticket. The particular network fault is entered into trouble field 62L and any additional data that may be helpful to the maintenance and repair personnel is entered in data field 62M. Field 62M may be subdivided into separate fields. Data

-11-

for trouble ticket 60 is automatically transmitted from network monitoring system 12 to fault processing system 18 via communication link 16.

The present invention, in fault management and resolution system 28 applies so-called "case-based reasoning" to trouble ticket resolution. The case-based reasoning method of the invention represents fault resolution expertise in the form of cases, i.e., particular, specific, fault resolution scenarios stored in completed trouble tickets, rather than general rules as are used in existing RBR systems. The present invention recalls, adapts (in one embodiment), and executes these prior fault resolution scenarios stored in prior trouble tickets in an attempt to resolve an outstanding trouble ticket. When confronted with a novel problem, fault resolution system 32, using case-based reasoning, retrieves one or more similar trouble tickets from trouble-ticketing system 34 and tries to adapt the resolution of the prior trouble ticket in an attempt to resolve the outstanding network fault. Each stored trouble ticket represents a prior case of problem solving, and the entire trouble ticket database stored in memory module 50 represents a case library.

FIG. 4 illustrates the overall processing steps carried out by fault management and resolution system 28 when processing a trouble ticket 60. In step 100, the system receives a trouble ticket from gateway 26, or information indicative of a network fault from user interface 38. Once a trouble ticket has been received, the system proceeds to step 102 in which, as will be explained in greater detail, fault resolution system 32 accesses trouble-ticketing system 34 to retrieve one or more stored trouble tickets from memory 50 that are similar to outstanding trouble ticket 60. From step 102, the fault resolution system proceeds to step 104 wherein the system selects and examines the retrieved trouble ticket which is most similar to outstanding trouble ticket 60. If the retrieved ticket is a perfect match in all relevant fields, the resolution is not altered and the system proceeds

-12-

directly to step 106. On the other hand, if the trouble ticket is not an exact match in all relevant fields, step 104 employs adaptation techniques, which will be described in more detail, to adapt the pre-existing resolutions to the present fault, and then proceeds to step 106.

In step 106, the system proposes potential solutions by displaying them on display monitor 46 or printing them on printer 48 to allow the maintenance and repair person to inspect, execute, and/or manually adapt the proposed solutions, if desired. The system may also test the potential resolutions in step 106 by transmitting them via communications link 36 to configuration management module 20 for execution by network control software. From step 106, the system proceeds to step 108 in which the new resolution correcting the network fault is processed. The resolution is entered into resolution field 62P of trouble ticket 60. Trouble ticket 60 is then stored in trouble ticket memory 50, thus adding to the system's knowledge base that may be accessed in order to resolve future communications network faults. In another embodiment, the resolution may be automatically transmitted via communication link 36 to configuration management module 20 to correct the outstanding network fault.

The system is capable of fully automatic operation and can receive relevant trouble tickets, process them, and correct network faults fully automatically without any user intervention.

FIG. 5 illustrates in more detail the operation of retrieve step 102. The usefulness of a trouble ticket retrieved from trouble ticket memory 50 depends upon the degree to which it is similar to outstanding ticket 60. Retrieve step 102 uses at least one set of determinators stored in determinator module 110 that includes relevance information defining relationships among classes of network faults and corresponding sets of trouble ticket data fields 62A-62R. A set of determinators includes at least one

-13-

determinator. A determinator identifies relevant trouble ticket data fields for specific kinds of troubles. For a given network fault, a set of determinators identifies the relevant fields of a trouble ticket for a given network fault so that only trouble tickets that are useful in resolving the outstanding trouble ticket are retrieved from trouble ticket memory 50.

As an example, the network fault recorded in data field 62L of trouble ticket 60 is "file_transfer_throughput = F". In order to select relevant trouble tickets from memory 50, the relevant data fields to be looked at are those that represent things such as bandwidth, network load, packet collision rate, and packet deferment rate. A determinator may thus be phrased as: "The solution to the trouble 'file transfer_throughput = F' is determined by looking at bandwidth, network load, packet collision rate, and packet deferment rate". A set of such determinators relating network faults to relevant data fields is provided and stored as data structures in module 110. When the system is processing an outstanding trouble ticket, it uses at least one of the sets of determinators stored in determinator module 110 and retrieves relevant trouble tickets from trouble ticket memory 50 as a function of the determinators.

FIG. 5 illustrates four different types of sets of determinators that may be stored in the system for use by retrieve step 102. Determinators may be stored as "macros" in macro data structure 112, "rules" in rules database data structure 114, as a decision tree in decision tree data structure 116, and/or as a neural network memory in neural network memory data structure 117. Macros may be determined by recording the actions of the repair and maintenance personnel during the sorting and selection of relevant trouble tickets and then defining the process as a series of steps that may be processed in step 102 when similar troubles arise in the future. The macro is a determinator that matches specific faults to trouble ticket data fields. Rules

-14-

for rules database 114 may be determined by having domain experts explicitly specify a set of rules that match specific faults to trouble ticket data fields. Each rule is a determinator. Using knowledge engineering techniques such as the "consult/implement/test" technique previously described, these rules can be refined manually, automatically, or by a combination of automatic and manual modification as the system deals with network faults, and can change as the network changes. A decision tree may be determined by applying an information theoretic induction algorithm such as the "ID3" algorithm to the trouble ticket database to provide a decision tree structure containing determinators that match specific faults to trouble ticket data fields. The "ID3" algorithm is described in "Learning Efficient Classification Procedures and their Applications to Chess End-Games" by J.R. Quinlan published in the book Machine Learning: An Artificial Intelligence Approach, pp. 463-482, 1983. A neural network memory may be determined by applying a neural network learning algorithm such as a "back-propagation" algorithm to the trouble ticket database to provide a neural network memory structure containing determinators that match specific faults to trouble ticket data fields. The "back-propagation" algorithm is described in "Learning Internal Representation by Error Propagation," by D.E. Rumelhart, G.E. Hinton, and R.J. Williams published in the book Parallel Distributed Processing: Exploration in the Microstructure of Cognition, MIT Press, 1986, pp. 318-362. These determinators may be used singularly or in combination in step 102 when selecting relevant trouble tickets.

Typically, when a group of trouble tickets has been retrieved, the tickets are indexed in order of increasing similarity to the outstanding network fault. The first so-indexed ticket is chosen and its resolution strategy is considered and/or executed by steps 104-108. If that resolution does not work, the next ticket is chosen and the process is repeated. An important feature of the present

-15-

invention is that as each retrieved trouble ticket is processed by step 108, a new trouble ticket can be created that records the system's experience with a given network fault and a given resolution and that new ticket placed in memory 50 for future reference. For example, if the system executed a resolution to the outstanding trouble ticket and the resolution proved to be unsatisfactory, the system can copy trouble ticket 60 to generate a new trouble ticket, store the unsatisfactory resolution in resolution field 62P, and then indicate in field 62Q that the resolution was unsatisfactory by filling in status "No Good." This new trouble ticket is then stored in memory 50, thus increasing the system's knowledge base so that the system learns, over time, which resolutions work and which do not for a particular network fault.

FIG. 6 illustrates in more detail the operation of adaptation step 104. In step 104, the resolution in a retrieved ticket can be adapted to resolve the outstanding network fault if no trouble ticket is located in memory 50 having a perfect match between the network fault, the relevant data fields, and the proposed resolution. FIG. 6 illustrates three adaptation techniques that the system can use to modify the resolution of a retrieved trouble ticket for use with an outstanding trouble ticket. Adaptation step 104 can access a data structure 118 containing a null adaptation technique, a data structure 120 containing a parameterized adaptation technique, and/or a data structure 122 containing an abstraction/respecialization adaptation technique.

The null adaptation technique provided by module 118 means that no adaptation is performed and the resolution in the retrieved trouble ticket is simply passed from step 104 to step 106. The system may be programmed by a user to provide null adaptation through user interface 38. This

-16-

technique allows the adaptation step to be bypassed and thereby allows the prior resolution to be filled in exactly in data field 62L.

The parameterized adaptation technique performed by module 120 is a method by which a solution variable in the representation of an outstanding network fault entered in data field 62L of an outstanding trouble ticket is adjusted relative to a problem variable in the representation of an outstanding network fault entered in data field 62L. The adjustment is based on the relationship between the solution and problem variables in the retrieved ticket. For example, assume that the network fault of a retrieved trouble ticket is "file_transfer_throughput = F" and the resolution of this fault is "A = f(F), adjust_network_load = A". If the outstanding trouble ticket fault is "file_transfer_throughput = F'", then the parameterized adaptation module 120 may propose resolution "A' = f(F')", adjust_network_load = A' ". Thus, parameterized adaptation module 120 chooses a resolution such that the relationship between F' and A' in the outstanding trouble ticket is the same as the relationship between F and A in the retrieved trouble ticket. Parameterized adaptation module 120 can provide the proposed resolution using a user-defined function that calculates A' for any values of F', a function expressed by a neural network or fuzzy logic system that calculates A' for any values of F', a lookup table, where the value of A' is determined by interpolation over values of F', a sequence of steps that calculates A' for any values of F', or a decision tree that calculates A' for any values of F'. The different types of parameterized adaptation steps may be used singularly or in combination.

The abstraction/respecialization technique provided by module 122 allows the system to consider viable alternative resolutions to a fault. If, for example, there is a prohibitive constraint on a proposed resolution, such as an unobtainable increase in bandwidth, the system abstracts over

-17-

i.e., stops considering and processing the retrieved trouble ticket that contains the proposed resolution and respecializes to i.e., chooses another retrieved trouble ticket that contains an alternative solution. For example, suppose the network fault is "file_transfer_throughput = F", and consider two retrieved tickets that are equally similar to the outstanding trouble ticket. The first ticket contains the following trouble (i.e., network fault) and resolution:

- 1) trouble: file_transfer_throughput = F
additional data: none
resolution: A = f(F), adjust_network_load = A
resolution status: good

The second trouble ticket contains the following trouble and resolution:

- 2) trouble: file_transfer_throughput = F
additional data: none
resolution: B = g(F), increase_bandwidth = B
resolution status: good

If there is a constraint on the possible resolution of the outstanding ticket such as a prohibition on the adjustment of network load or, if the execution of "adjust_network_load" does not resolve the trouble, the system would proceed from proposing the resolution of the first retrieved trouble ticket to proposing the resolution of "increase_bandwidth" of the second retrieved trouble ticket. A new trouble ticket is then created and stored in trouble ticket memory 50 containing the following information:

- 3) trouble: file_transfer_throughput = F
additional data: adjust_network_load = no
resolution: A = f(F), adjust_network_load = A
resolution status: no good

The system creates a new trouble ticket that may be stored in trouble ticket memory 50 containing the trouble, proposed

-18-

resolution, and constraint for use in resolving future network faults, thus increasing the system's knowledge base.

FIG. 6A illustrates an additional adaptation technique provided by critic-base adaptation module 124. Critic-based adaptation module 124 allows a user, through user interface 38 to edit a proposed displayed potential solution presented by propose step 106 or to enter his or her own solution to the outstanding network fault. Critic-based adaptation is another form of adaptation that allows the system to adapt previous resolutions to novel network faults. Critic-based adaptation includes adding, removing, reordering, or replacing steps in the proposed retrieved solution. For example, considering the first retrieved trouble ticket described in connection with FIG. 6 above, a maintenance and repair person could include the data field "network_load" and refine the solution by providing a two-place function $f(F,N)$ that calculates the amount of adjustment based on the values of file "transfer_throughput" and "network_load". The proposed resolution to the new trouble would be the following:

- 4) trouble: file_transfer_throughput = F
additional data: network_load = N
resolution: $A = f(F,N)$, adjust_network_load = A
resolution status: good

This new trouble ticket is processed in step 108 and stored in trouble ticket memory 50. If the same network fault recurs and the value of network load is available, the system in step 104 could propose a value for A using the two-place function $f(F,N)$. If the value of network load is unavailable, the value of A would be determined using the one place function $f(F)$. Note also that parameterized adaptation of module 120 could be applied to this new trouble ticket for use in resolving a future network fault.

The system is capable of using null adaptation, parameterized adaptation, abstraction/respecialization, and critic-based adaptation singularly or in combination. The use of adaptation techniques allows the system to increase

-19-

its knowledge base and learn from prior fault resolution scenarios in order to more efficiently deal with novel network faults.

FIG. 7 is a pseudo-code implementation of the processing steps illustrated in FIGS. 4-6A. FIG. 7 represents one pass through the system operation. FIG. 7 illustrates one determinator rule being selected and the use of parameterized adaptation to adapt the retrieved ticket. To enable the abstraction/respecialization adaptation technique and to be able to retrieve multiple tickets, the operations Retrieve, Adapt, and Propose are enclosed in a loop controlled by the command "Repeat until no other D_i is applicable". To provide for critic-based adaptation, the Propose module would allow the user to overwrite $V_{n+1}(f_r)$. To disable the adaptation operation, the code in the Adapt module is replaced with Set $V_{n+1}(f_r) = V_i(f_r)$ in which case the solution in ticket T_i is written directly to the outstanding ticket T_{n+1} without alteration.

Having thus described one particular embodiment of the invention, various alterations, modifications, and improvements will readily occur to those skilled in the art. For example, one skilled in the art will recognize that the present invention is applicable to networks other than local area networks. In addition, the present invention can be used with communications networks fault management systems other than trouble ticket type systems, such as spread sheet systems or database systems. Such alterations, modifications, and improvements are intended to be part of this disclosure and are intended to be within the spirit and scope of the invention. Accordingly, the foregoing description is by way of example only and is not intended as limiting. The invention is limited only as defined in the following claims and equivalents thereto.

What is claimed is:

-20-

CLAIMS

1. A method of resolving faults in a communications network, comprising the steps of:

acquiring information about an outstanding communications network fault;

retrieving relevant, previously stored sets of data regarding prior communications network faults and resolutions of those prior communications network faults from a database;

using at least a portion of one of the retrieved sets of data to provide a resolution of the outstanding communications network fault.

2. The method of claim 1, wherein the sets of data include fault and resolution data.

3. The method of claim 1, further comprising the step of using at least one determinator to correlate the outstanding communications network fault to data fields in the sets of data and wherein the retrieving step retrieves sets of data as a function of the at least one determinator.

4. The method of claim 3, further comprising the steps of determining the at least one determinator by recording problem solving actions of at least one communications network user using a trouble ticket data structure and using the problem solving actions to define a series of steps that correlate the outstanding communications network fault to data fields in a set of data fields in the trouble ticket data structure.

5. The method of claim 3, further comprising the step of determining the at least one determinator by determining a set of rules that matches the outstanding communications network fault to a set of data fields in a trouble ticket data structure.

-21-

6. The method of claim 3, further comprising the step of determining the at least one determinator by applying an information theoretic induction algorithm to a set of trouble tickets stored in a trouble ticket library.
7. The method of claim 6, wherein the step of applying an information theoretic induction algorithm includes applying an ID3 algorithm.
8. The method of claim 3, further comprising the step of determining the at least one determinator by applying a neural network learning algorithm to a set of trouble tickets stored in a trouble ticket library.
9. The method of claim 8, wherein the step of applying a neural network learning algorithm includes applying a back propagation algorithm.
10. The method of claim 1, wherein the step of using at least one determinator includes the step of using a set of determinators to correlate the outstanding communications network fault to data fields in a set of data fields of a trouble ticket data structure.
11. The method of claim 1, wherein the step of using at least a portion of the resolution of one of the retrieved sets of data to provide a resolution of the outstanding communications network fault further comprises the step of adapting the resolution from at least one similar retrieved trouble ticket to provide the resolution of an outstanding trouble ticket.

-22-

12. The method of claim 11, wherein the step of adapting includes using a null adaptation technique wherein a resolution of a retrieved trouble ticket is stored in a data field in the outstanding trouble ticket capable of storing the resolution.

13. The method of claim 11, wherein the step of adapting includes using a parameterized adaptation technique including the steps of determining a relationship between a data field and a resolution in a retrieved trouble ticket and providing a resolution of the outstanding trouble ticket using the relationship.

14. The method of claim 11, wherein the step of adapting includes using an abstraction/respecialization technique including the steps of examining a retrieved trouble ticket for a constraint on a resolution stored in the retrieved trouble ticket, determining if the constraint applies to the outstanding trouble ticket, and proceeding to examine another retrieved trouble ticket if the constraint applies to the outstanding trouble ticket.

15. The method of claim 11, wherein the step of adapting includes modifying, by a communications network user, a resolution in a retrieved trouble ticket.

16. The method of claim 11, further comprising the step of proposing a potential resolution prior to providing a resolution of the outstanding trouble ticket.

17. The method of claim 16, wherein the step of proposing potential resolutions includes the step of displaying potential resolutions to a communications network user.

-23-

18. The method of claim 16, wherein the step of proposing potential resolutions includes the step of printing potential resolutions on a printer.
19. The method of claim 16, wherein the step of proposing potential resolutions includes the step of automatically transmitting a proposed resolution to a network configuration module for automatic execution of the proposed resolution.
20. . The method of claim 16, further comprising the step of storing the outstanding trouble ticket in the trouble ticket library after an appropriate resolution has been determined and stored in the data field capable of storing the resolution.
21. The method of claim 16, further comprising the steps of storing copies of the outstanding trouble ticket in the trouble ticket library, each copy having a potential resolution stored in the data field capable of storing the resolution and an indication in another data field of whether or not the proposed resolution corrected the network fault.
22. An apparatus for resolving faults in a communications network, comprising:
 - an input system acquiring information about an outstanding communications network fault in the form of an outstanding trouble ticket;
 - a fault management system storing trouble tickets containing data regarding prior communications network faults and resolutions of those prior communications network faults in a trouble ticket library;

-24-

a fault resolution system retrieving a set of trouble tickets that are similar to the outstanding trouble ticket and providing a resolution of the outstanding trouble ticket based on resolutions contained in the retrieved set of trouble tickets; and

a system transmitting the resolution of the outstanding trouble ticket to the communications network.

23. The apparatus of claim 22, wherein the fault resolution system further comprises a system, coupled to the adaptation system, proposing potential resolutions of the outstanding trouble ticket.

24. The apparatus of claim 23, further comprising a user interface, coupled to the system proposing potential resolutions, allowing editing of potential resolutions.

25. A method of resolving faults in a communications network, comprising the steps of:

providing a trouble ticket data structure capable of describing information concerning a communications network fault;

using the trouble ticket data structure to create a trouble ticket that describes a specific network fault, the trouble ticket having at least one data field capable of storing information regarding the network fault and a data field capable of storing a resolution of the specific network fault;

providing a trouble ticket library capable of storing at least one trouble ticket;

receiving an outstanding trouble ticket that describes an outstanding communications network fault;


using at least one determinator to correlate the outstanding communications network fault to data fields in the set of data fields of the trouble ticket data structure;

-25-

retrieving a set of trouble tickets from the trouble ticket library that are similar to the outstanding trouble ticket as a function of the at least one determinator; and using at least a portion of the resolution from at least one trouble ticket to provide a resolution of the outstanding trouble ticket.


1/7

Fig.1



2/7

Fig.2



3/7


60

Fig. 3

624	Entry-ID	000000000116	Submitter	SPECTRUM	62G
628	Alarm ID	57	Create-date	07/24/92 08:51:08	62H
62C	Alarm Date/Time	07/24/92 08:51:07	Notify-method	o None o Notifier o E-mail	62T
62S	Condition	o Red o Orange o Yellow	Assigned-Priority	o Low o Medium o High	62U
62D	Device Name	Ethernet-Randy	Assigned-to		62I
62E	Device Type	Ethernet	Last-modified-by	SPECTRUM	62J
62F	IP Address		Modified-date	07/24/92 08:51:08	62K
Trouble			file_transfer_throughput=F		62L
Additional Data			network_load=20,collision_rate=15,deferment_rate=20,users=31		62M
History of Trouble					62N
Probable Cause					62O
Resolution					62P
62Q	Resolution Status	o Good o No Good o InProgress			
62R	Ticket Status	o New o Assigned o Rejected o Closed			


4/7

Fig.4



5/7

Fig.5



6/7

Fig.6





Fig.6A




7/7

Fig. 7

Given:
 A set of tickets $T = \{T_1, T_2, \dots, T_n\}$ and a set of ticket fields $F = \{f_1, f_2, \dots, f_n\}$ such that each ticket T_i is of the form $\{V_i(f_1), V_i(f_2), \dots, V_i(f_n)\}$ where f_1, \dots, f_n are background description fields, f_p is a problem description field, f_r is a resolution description field, and $V_i(f_n)$ is the value of field f_n in ticket T_i .

A set of determinator/adaptation rules $D = \{D_1, D_2, \dots, D_m\}$ such that each D_i is of the form:
 If $V(f_p)$ then the set of relevant fields is $R = \{f_1, f_2, \dots, f_m\}$ and $V(f_r) = f(V(f_1), V(f_2), \dots, V(f_m))$,
 where $R \subseteq F$ and $V(f_r)$ is determined by some function f over the values of the fields in R .

Acquire:
 Enter new trouble ticket T_{n+1} where $V_{n+1}(f_r) = \text{null}$



INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 94/02056

A. CLASSIFICATION OF SUBJECT MATTER
IPC 5 H04L12/24 H04L12/26 H04M3/24 G06F11/22

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 5 H04L G06F H04M

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P, X	<p>The Ninth Conference on Artificial Intelligence for Applications, 01.-05.03.1993, Orlando, US, pages 114-120, IEEE, New York, US; XP379594;</p> <p>L. LEWIS: 'A Case-Based Reasoning Approach to the Management of Faults in Communications Networks.' see page 115, right column, paragraph 2 - page 119, right column, paragraph 2; figures</p> <p>---</p> <p>-/-</p>	1-25

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *'A' document defining the general state of the art which is not considered to be of particular relevance
- *'E' earlier document but published on or after the international filing date
- *'L' document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *'O' document referring to an oral disclosure, use, exhibition or other means
- *'P' document published prior to the international filing date but later than the priority date claimed

- *'T' later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *'X' document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *'Y' document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *'&' document member of the same patent family

Date of the actual completion of the international search

21 June 1994

Date of mailing of the international search report

30.06.94

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax (+31-70) 340-3016

Authorized officer

Gries, T

INTERNATIONAL SEARCH REPORT

Inte onal Application No
PCT/US 94/02056

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P,X	IEEE NETWORK: THE MAGAZINE OF COMPUTER COMMUNICATIONS, vol.7, no.6, November 1993, NEW YORK US pages 44 - 51 L. LEWIS/G. DREO: 'Extending Trouble Ticket Systems to Fault Diagnostics.' see page 47, left column, paragraph 5 - page 48, right column, paragraph 1 ---	1-25
A	GLOBECOM '90. IEEE Global Telecommunications Conference & Exhibition, 02.-05.12.1990, San Diego, US, vol. 2, pages 1344-1350, IEEE, New York, US; XP221102; Y. LIROV/O.-C. YUE: 'Expert maintenance systems in telecommunications.' see page 1345, right column, line 21 - line 40 ---	1-3,6
A	EP,A,0 442 809 (DIGITAL EQUIPMENT CORPORATION) 21 August 1991 see abstract see page 2, line 31 - page 3, line 15; claims ---	1-3,6
A	AT & T TECHNICAL JOURNAL, vol.67, no.6, November 1988, NEW YORK US pages 137 - 154, XP212378 T.E. MARQUES: 'StarKeeper network troubleshooter: an expert system product.' see page 139, left column, paragraph 2 -paragraph 3; figure 2 see figure 6 see page 147, right column, paragraph 5 - page 149, right column, paragraph 2; figure 7 ---	1-4,6, 22,25
A	EP,A,0 508 571 (HEWLETT-PACKARD COMPANY) 14 October 1992 see abstract see page 3, line 25 - line 49 see page 12, line 13 - line 26; claims -----	1-3,6

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/US 94/02056

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
EP-A-0442809	21-08-91	JP-A-	5313940	26-11-93
EP-A-0508571	14-10-92	JP-A-	5083257	02-04-93